JUNE 12, 2020 BY JASON K. RESCH

When will AI take over?

Watch Episode Download Article

Artificial Intelligence has achieved extraordinary feats. Given the current state and trajectory of AI, for how much longer can humans remain the *most intelligent creature* on Earth?

We’re not the strongest or hardiest. Nor are we the fastest reproducing. Yet humans dominate life on Earth.

By some estimates we’ve co-opted 40% of the Earth’s terrestrial photosynthetic capacity. Humans inhabit every continent and maintain a permanent presence under the ocean and in outer space. We’re the only species to have stepped foot on other worlds or to have released the energy of the atom.

Our success as a species we owe to our intelligence--an intelligence that lets us coordinate actions, share information, and build tools. But what will happen when we finally build something smarter than us?

In just the last decade, we’ve witnessed the rise of AI that can hold conversations, speak any of 100 languages, beat anyone at Jeopardy!, Chess, Go, Poker, and in any Atari video game. AI recognizes objects and faces better than us—even spots cancer better than most doctors. AI has been used to invent technology, discover laws of physics, even identify new drugs. Today’s AI can compose as well as Bach, and paint as well as van Gogh. It has developed its own highly original artistic styles. It can drive, fly, and even move better than most humans.

What chance do we have in the coming age where AI eclipses humanity in all areas?

This article will explore that question. We will review the areas where AI has already achieved superhuman abilities, and cover the dwindling areas where humans still hold an edge.

It is a most-pressing question. Though we have held the title of “most

SEARCH

Search …

BE NOTIFIED OF NEW POSTS

Name*

Email*

SUBSCRIBE

FOLLOW ALWAYS ASKING

POPULAR ARTICLES

Is there life after death?
251.2k views | posted on August 24, 2020

What is the meaning of life?
75.8k views | posted on November 30, 2020

Are we alone?
52.7k views | posted on May 26, 2020

What is time?
43.1k views | posted on August 6, 2020

How big is the universe?
18.9k views | posted on May 18, 2020

Why does anything exist?
18.3k views | posted on March 8, 2021

When will AI take over?
11.5k views | posted on June 12, 2020

intelligent species” for a hundred thousand years, we could be relieved of this title in a few decades.

We must prepare for this transition.

Contents

1. What is Intelligence?
    1.1. Narrow vs Strong AI
2. Human Intelligence
    2.1. Power of the Human Brain
    2.2. Abilities of the Human Brain
3. The Present State of Artificial Intelligence
    3.1. Communication abilities of AI
        3.1.1. Transcription
        3.1.2. Translation
        3.1.3. Comprehension
        3.1.4. Putting it all together
    3.2. Learning abilities of AI
        3.2.1. Video Games
        3.2.2. Board Games
        3.2.3. Partial Information Games
        3.2.4. Real-time Games
        3.2.5. Beyond Games
    3.3. Movement abilities of AI
        3.3.1. Driving
        3.3.2. Flying
        3.3.3. Nimble Robots
    3.4. Recognition abilities of AI
    3.5. Creative abilities of AI
        3.5.1. Art
        3.5.2. Music
        3.5.3. Imagination
        3.5.4. Invention
    3.6. Reasoning abilities of AI
4. Artificial Intelligence and Jobs
    4.1. Jobs Taken by AI and Automation
    4.2. Jobs Threatened by AI
        4.2.1. AI in Law
        4.2.2. AI in Medicine
        4.2.3. AI in Science
        4.2.4. AI in Finance
        4.2.5. AI in Entertainment
        4.2.6. AI in Transportation
        4.2.7. AI in Writing
    4.3. What Jobs are Safe?
5. The Future of Artificial Intelligence
    5.1. Algorithms of Intelligence
        5.1.1. Neural Networks
    5.2. Computing Trends
    5.3. An Intelligence Explosion
        5.3.1. The Doomsday Equation
        5.3.2. A Singularity in History
    5.4. Superpowers of Superintelligence
    5.5. Limits of Intelligence
6. Conclusions
    6.1. Are Intelligent Machines Possible?
    6.2. What does our Future Hold?
    6.3. When will AI take over?

What is Intelligence?

Before we can discuss artificial intelligence, we need to start with an

Can aging be cured?
7.6k views | posted on July 30, 2020

Is the universe fine-tuned?
3.5k views | posted on October 14, 2020

LATEST ARTICLES

Why does anything exist? March 8, 2021

What is the meaning of life? November 30, 2020

Is the universe fine-tuned? October 14, 2020

Is there life after death? August 24, 2020

What is time? August 6, 2020

RECENT COMMENTS

Guilherme Monteiro Jr. on How big is the universe?

Guilherme Monteiro Jr. on How big is the universe?

Guilherme Monteiro Jr. on How big is the universe?

Guilherme Monteiro Jr. on How big is the universe?

Guilherme Monteiro Jr. on How big is the universe?

Guilherme Monteiro Jr. on How big is the universe?

Guilherme Monteiro Jr. on How big is the universe?

Guilherme Monteiro Jr. on How big is the universe?

Guilherme Monteiro Jr. on How big is the universe?

Guilherme Monteiro Jr. on How big is the universe?

REGISTER TO COMMENT

Sign Up

Login

agreement on what *intelligence* is.

According to the *agent-environment interaction* model of intelligence, something is intelligent if it:

> "perceives its environment and interacts with it in a manner consistent with achieving a goal."

This definition captures the full spectrum of intelligent behavior, regardless of how simple or complex it is. It includes creatures from worms to humans, and machines from thermostats to chess playing AIs.

In the agent-environment interaction model of intelligence, an intelligent agent acts upon the environment, the environment changes, the agent perceives these changes and the agent can choose a new course of action.

Anything fitting this definition of intelligence is an *intelligent agent*.

Narrow vs Strong AI

Only with a goal in mind can we tell unintelligent behaviors from intelligent ones. Accordingly, every intelligent agent, be it an animal, robot, AI, extraterrestrial, or otherwise, must have goals to behave intelligently.

Agent	Goal	Environment	Actions	Perceptions
Thermostat	Maintain warmth	Room	Turn on heat	Temperature reading
Chess AI	Win game	Chess board	Move piece	Opponent's response
Delivery drone	Deliver Package	Airspace	Adjust rotors	Air Speed, GPS coordinates
Human being	Survive & Thrive	The world	Control muscles	Sights, sounds, smells, etc.

Every intelligent agent, from a chess playing AI to a human being can be framed in the terms of its: ***goals***, ***environment***, ***actions***, and ***perceptions***.

By this definition, we have achieved artificial intelligence. But existing artificial intelligences have narrowly defined goals and are usually competent only in one specific function.

AIs that are only good at one thing, be it playing chess or keeping an airplane level are known as narrow AI. This is in contrast to *Strong AI*, an AI with wide-ranging and general purpose intelligence.

Definition: **Strong AI** – an AI that can perform any mental task a human is able to perform.

While *narrow AI* can outperform humans in many tasks, human intelligence is general.

In order to estimate how far we are from *Strong AI*, we must consider all the aspects of human intelligence and see how near or far AI is from achieving parity. Only when AI meets or exceeds human ability in each dimension of intelligence will Strong AI be achieved.

Human Intelligence

In order for AI to take over, it is going to have to meet and exceed human intelligence. How far off is present computing technology from matching the power of the human brain?

Power of the Human Brain

Human Intelligence involves interactions between 100 billion neurons linked by 1,000 trillion connections.
There are as many neurons in your brain as stars in the Milky Way Galaxy.

The human brain has about 100 billion neurons, roughly the number of stars in our galaxy.

Each neuron has about 10 thousand synapses — connections to other neurons. This gives a total of one quadrillion (10^{15}) connections in the human brain.

Each of the neurons in the human brain can *fire*--signal to other neurons--at up to 1,000 times per second. The 10^{15} synapses signaling at 1,000 times per second gives a synaptic signaling rate of 10^{18} per second.

Every second, your brain sends as many synaptic signals as there are grains of sand on all of Earth's beaches – (estimated to be on the order of $10^{18} = \text{1,000,000,000,000,000,000}$).

This rate of processing, 10^{18} operations per second, is known as the *exaop scale*. Exa- is the SI prefix for 10^{18} and *op* is short for operation.

In 2018, the world's fastest supercomputer broke the exaop scale.

This supercomputer is called *Summit*. It was developed by IBM and is operated by the U.S. Department of Energy at the Oak Ridge National Laboratory in Tennessee.

Summit is the world's fastest super computer. Image Credit: Oak Ridge National Laboratory

To reach the processing rate of the human brain, Summit requires 2.4 million CPUs, 10 million watts of power, 340 tons of equipment, and two tennis courts of space.

The human brain performs as many operations per second as Summit, but uses just 20 watts of power, weighs just 1.5 kilograms and fits in a two-liter bottle with room to spare.

While our computing speed may have finally caught up with our biology, it has a long way to go to meet the efficiency and compactness of the human brain. Moreover, human intelligence requires more than raw computing power. It needs the right algorithms, circuitry, and programming.

All the computers in the world won't yield general purpose AI without *the right software*.

Abilities of the Human Brain

Human intelligence is multifaceted. It embodies many abilities such as: problem solving, reasoning, pattern recognition, creativity, learning, language, planning, intuition, and applying knowledge.

If artificial intelligence is to ever achieve human intelligence, it must succeed in all these areas.

As of 2020, artificial general intelligence has not been achieved. But that isn't stopping anyone from trying. A 2017 survey found 45 R&D projects by different organizations, all of them pursuing strong AI.

> Each practitioner thinks there's one magic way to get a machine to be smart, and so they're all wasting their time in a sense. On the other hand, each of them is improving some particular method, so maybe someday in the near future, or maybe it's two generations away, someone else will come around and say, 'Let's put all these together,' and then it will be smart.
>
> *— Marvin Minsky, co-founder of the MIT AI laboratory*

If Minsky is right, general intelligence is what you get from combining a bunch of distinct abilities together. Accordingly, human general intelligence can be viewed as the combined abilities to:

- Communicate via natural language
- Learn, adapt, and grow
- Move through a dynamic environment
- Recognize sights and sounds
- Be creative in music, art, writing and invention
- Reason with logic and rationality to solve problems

Where do we stand in terms of progress in each of these areas?

In recent years, extraordinary achievements have been made in each of these domains. Be warned: if you're not familiar with recent progress in AI the next section will be shocking.

The Present State of Artificial Intelligence

It is one thing to guess what might be possible in the future. It's another to confront the hard facts of what's already been done.

Communication abilities of AI

One ability of human intelligence is nothing short of telepathy. Thoughts in one brain can spread through vibrations in the air to generate similar thoughts in other nearby brains.

This is the magic of language. Through writing, this power is further enhanced to allow brains to share ideas with other humans thousands of years in the future or on the other side of the globe.

The scribbles on your screen are performing this feat of magic right now.

Language involves the separate functions of: transcription (turning sounds into words), translation, comprehension, and to give a response, speech synthesis (turning words into sounds).

Transcription

Voice transcription, also known as speech recognition, converts spoken words to text.

Early transcription software has existed since the 1980s, but it was primitive by today's standards. The first versions required voice training for each speaker, had limited vocabularies, and speakers had to pause between each word.

But in 2016, Microsoft's Artificial Intelligence and Research Unit built the first speech recognition technology to surpass the accuracy of human transcriptionists.

Humans, working on a set of recordings had an error rate of 5.9%. By a narrow margin, Microsoft's system achieved a lower error rate.

As of 2020, the latest commercial version of Dragon NaturallySpeaking supports speech at up to 160 words per minute, at an accuracy of 99% without training.

Translation

Modern AI has an incredible capacity for language translation.

Google Translate supports over 100 languages, and it is entirely self-taught.

Rather than hand program language translation rules, Google engineers built an AI that could teach itself. It is called *Google Neural Machine Translation* (GNMT).

Google Translate can perform on-the-fly translations of signs and other written text. Image

Credit: Google

Google Engineers fed GNMT example translations published by the United Nations and the European Parliament. Just as humans learned to decipher ancient Egyptian from the example on the Rosetta Stone, GNMT figured out how to translate between 101 human languages.

The top human polyglot in the world *only* knows 59 languages.

Comprehension

Comprehension is the ability to extract the meaning conveyed by words.

In 2019, engineers at OpenAI succeeded to such a degree that they were scared by their creation. They considered their AI so dangerous they said it would be irresponsible to release to the public.

These engineers created an AI called *Generative Pretrained Transformer 2* or GPT-2.

The AI had a simple task: given an input text, predict the most likely next word to follow.

After training on 8 billion documents (comprising 40 gigabytes) of English text, GPT-2 showed surprising abilities: GPT-2 could invent pages of coherent text. Moreover, GPT-2 demonstrated reading comprehension. It could even summarize and answer questions about an input text — despite not being trained for that!

Given any prompt, GPT-2 can generate as much or as little additional text as desired. For example, when given the word "**Today**", GPT-2 continued with "**Today, the United Nations has called for the immediate withdrawal of all nuclear weapons from the world.**"

GPT-2 could even write whole articles, complete with fake quotes from scientists. The AI demonstrated an understanding of world geography and even naming conventions used in different parts of the world.

Researchers fed GPT-2 the prompt: "**In a shocking finding, scientist discovered a herd of unicorns living in a remote, previously unexplored valley, in the Andes Mountains. Even more surprising to the researchers was the fact that the unicorns spoke perfect English.**"

They asked GPT-2 to continue from there. The following is the text the AI came up with:

> The scientist named the population, after their distinctive horn, Ovid's Unicorn. These four-horned, silver-white unicorns were previously unknown to science.
> Now, after almost two centuries, the mystery of what sparked this odd phenomenon is finally solved.
> Dr. Jorge Pérez, an evolutionary biologist from the University of La Paz, and several companions, were exploring the Andes Mountains when they found a small valley, with no other animals or humans. Pérez noticed that the valley had what appeared to be a natural fountain, surrounded by two peaks of rock and silver snow.
> Pérez and the others then ventured further into the valley. "By the time we reached the top of one peak, the water looked blue, with some crystals on top," said Pérez.
> Pérez and his friends were astonished to see the unicorn herd. These creatures could be seen from the air without having to move too much to see them – they were so close they could touch their horns.
> While examining these bizarre creatures the scientists discovered that the creatures also spoke some fairly regular English. Pérez stated, "We can see, for example, that they have a common 'language,' something like a dialect or dialectic."
> Dr. Pérez believes that the unicorns may have originated in Argentina, where the animals were believed to be descendants of a lost race of people who lived there before the arrival of humans in those parts of South America.
> While their origins are still unclear, some believe that perhaps the creatures were created when a human and a unicorn met each other in a time before human civilization. According to Pérez, "In South America, such incidents seem to be quite common."
> However, Pérez also pointed out that it is likely that the only way of knowing for sure if unicorns are indeed the descendants of a lost alien race is through DNA. "But they seem to be able to communicate in English quite well, which I believe is a sign of evolution, or at least a change in social organization," said the scientist.

While the story is not perfect, GPT-2 connected the prompt of the Andes mountains with Argentina, the University of La Paz, and a Spanish name Jorge Pérez, all on its own. The AI also chose Pérez to be an evolutionary biologist — exactly the field you might expect for someone who discovers a species.

OpenAI thought a system that could generate fake stories had potential for abuse.

> Due to our concerns about malicious applications of the technology, we are not releasing the trained model. As an experiment in responsible disclosure, we are instead releasing a much smaller model for researchers to experiment with.
>
> *— OpenAI*

The creators feared their technology would be exploited to create fake news or generate spam.

Tristan Greene said, “I’m terrified of GPT-2 because it represents the kind of technology that evil humans are going to use to manipulate the population – and in my opinion that makes it more dangerous than any gun.”

After nine months of debate between OpenAI and the broader community, OpenAI reversed its decision. GPT-2 is now available to the public for anyone to download and use.

Anyone can experiment with it online at: talktotransformer.com.

Putting it all together

In a stunning example, engineers behind *Google Assistant* put all these technologies together. The result is an AI called *Duplex*. Duplex can call salons and restaurants to make reservations all by itself.

Google Assistant will soon be able to call restaurants and make a reservation for you

In 2018, Google Assistant combined many facets of natural language and could schedule appointments over the phone.

Duplex was told to “arrange a haircut appointment on Tuesday morning anytime between 10 and 12.”

Duplex called a local hair salon and had the following conversation with the person at the other end of the line:

> (Ring)
> Person: Hello, how can I help you?
> **Duplex: Hi, I'm calling to book a women's haircut for a client. Umm.. I'm looking for something on May 3rd.**
> Person: Sure, give me one second.
> **Duplex: Mm-hmm.**
> Person: Sure, what time are you looking for around?
> **Duplex: At 12:00 PM.**
> Person: We do not have a 12:00 PM available. The closest we have to that is a 1:15.
> **Duplex: Do you have anything between 10:00 AM, and uhh, 12:00 PM?**
> Person: Depending on what service she would like. What service is she looking for?
> **Duplex: Just a woman's haircut, for now.**
> Person: Okay, we have a 10 o'clock.
> **Duplex: 10:00 AM is fine.**
> Person: Okay what's her first name?
> **Duplex: The first name is Lisa.**
> Person: Okay perfect. So I will see Lisa at 10 o'clock on May 3rd.
> **Duplex: Okay great, thanks.**
> Person: Great. Have a great day. Bye.

Duplex combined many language technologies, including transcription, comprehension, and speech synthesis. The result was a call where the person had no idea they were speaking to an artificial intelligence.

We might consider this a pass of the *Turing test*.

Learning abilities of AI

DeepMind was founded in 2010 with the mission of "understanding and recreating intelligence itself."

By 2013, DeepMind had made a general purpose learning algorithm. To prove it, they set it up to play old Atari video games — games it had never seen before, nor been told how to play.

Video Games

The AI could see the screen and it had a controller. The only thing it was told to do was to try to get the highest score it could. At first, it played like a child might play, randomly pressing the controls and quickly losing.

But after a short time it improved. After 10 minutes playing the paddle game *Breakout*, the AI could competently hit the ball several times before missing.

After two hours of playing, it had learned to play the game as well as the best human players.

The engineers decided to let the system keep running overnight.

Is Space Invaders the future of artificial intelligence? | Channel 4 News

A single learning algorithm learns to master a wide variety of computer games just by playing them.

When they checked the progress several hours later they were astonished — the system discovered new strategies for winning in some of the games. Of the seven games tested, their AI beat all previous AIs in six of the games, and learned to play better than any human player in three of the games.

This accomplishment so impressed Google that they decided to buy DeepMind for over $500 million in 2014.

DeepMind continued working on their learning algorithm. By 2015, their system could learn how to play over half of 49 tested games at a superhuman level. But its performance still lagged behind human players in some of the games. Something was holding the AI back.

By 2020, the engineers discovered the missing ingredient: *curiosity*. After adding curiosity to the learning algorithm, the newest AI called *Agent57*, learned how to play all 57 Atari games better than any human.

After DeepMind's success in Atari games they turned their attention toward the ancient board game Go.

Board Games

Go is not only the oldest continually played board game, but the most complex. The number of possible moves per turn is so great that all previous attempts at building a strong Go AI failed.

DeepMind engineers thought their learning algorithm might be succeed where others failed. They designed a Go-playing AI that learned by playing itself. They called it *AlphaGo*.

By 2016, DeepMind estimated AlphaGo to be strong enough to beat the best human players in the world — a feat that had never been done. To put it to the test, Google organized a match between AlphaGo and the world Go champion Lee Sedol. Go is popular in Asia and the event was highly televised.

More than 200 million people tuned in — twice the viewership of the Super Bowl.

World champion go player Lee Sedol moments before resigning the 5th game.

Despite going into the match with high confidence, AlphaGo beat Lee Sedol, in 4 of the 5 games.

One year later, DeepMind made a general purpose AI for learning to play board games given no information aside from the rules of the game. Because it started with zero knowledge, they called it *AlphaZero*.

AlphaZero could learn any game given the rules. When provided the rules of chess, AlphaZero learned to play the game better than any human in just four hours! In the time between breakfast and lunch, AlphaZero rediscovered all the common openings that took human chess masters centuries to work out.

AlphaZero had not only learned to play chess at a higher level than any human, it was better than the best human-programmed chess systems.

You could say AlphaZero didn't just learn how to play Chess better than any human — it learned how to program a computer to play chess better than any team of human programmers.

> I always wondered how it would be if a superior species landed on earth and showed us how they played chess. Now I know.
>
> *— Peter Heine Nielsen, Chess Grandmaster, on AlphaZero's chess style*

Because it learned everything from scratch, AlphaZero's playing style was unique. DeepMind founder Demis Hassabis, himself a former chess prodigy, said "It's like chess from another dimension."

In 2018, after 30 hours of self-play, AlphaZero learned how to play Go better than AlphaGo, and beat it 100-0.

This development was too much for Lee Sedol. In 2019, he announced his retirement from the game, saying "Even if I become the number one, there is an entity that cannot be defeated."

Partial Information Games

Board games are a far cry from the dynamic and unpredictable world we find ourselves in. In chess and go, the player has complete knowledge about everything. The entire board is always visible. There are no unknowns.

In the real world, however, we have to make do with partial and incomplete information. Given our uncertainties, we must act on guesses, intuition, and faith.

AI had proven itself in games with complete information. Could it also succeed in games more like the real world? Games with many actors, unknown variables, and constant change?

Could machines develop *intuition*?

A joint team between Facebook AI and Carnegie Mellon University aimed to find out.

They chose a game where uncertainty, partial information, and intuition reign supreme: *Poker*.

In Poker, players know their own cards, but can only guess what their opponents have. Image Credit: EVGENY KOVALEV SPB/SHUTTERSTOCK

They wanted to build the first AI that could beat professional poker players in the most popular--and most difficult--version of the game: *Texas Hold'em* with multiple players and no betting limit.

Previous AIs had success in one-on-one games, or in *limit poker*, but no limit poker and having multiple simultaneous opponents makes the game significantly more complex. No one had succeeded here.

In poker, players know only some of the cards. Cards held by opponents are not seen until the end. Deception becomes a necessary strategy by *bluffing* — pretending to have a stronger hand than one really holds.

The team adopted the strategy of AlphaZero. Their AI, called *Pluribus*, would teach itself the game by playing against itself, with no outside instruction.

In seven hours Pluribus learned to play Poker as well as the average person. After 20 hours it reached the level of professional human players. The researchers let it keep going. After 8 days, the team believed Pluribus had achieved definitively superhuman levels of Poker.

To prove it, they organized a tournament between Pluribus and professional human poker players. Each human player was among the best in the world — all had won over $1 million through playing poker.

Pluribus won decisively. It had learned more about Poker in a week than humans found in a century of playing.

> Pluribus is a very hard opponent to play against. It's really hard to pin him down on any kind of hand.
>
> *— Chris Ferguson, six-time World Series of Poker winner, who lost to Pluribus*

Noam Brown, of Facebook AI said of Pluribus, "It can bluff better than any human."

Real-time Games

Though poker is like the real world in that it includes hidden information. It is still like board games in that it is turn-based. In a turn-based game, players need only decide *how* to act, but not *when* to act.

But the real world is dynamic and change is constant. We have to decide not only how to act, but when to act. We also don't have immediate knowledge about what others are doing or have done.

This property of "real-time" is wholly absent from games like Chess, Go and Poker, but it is a core element of other game genres, such as *Real-time Strategy* games and *First-person Shooters*.

In a surprise showing at Valve's $24 million video game tournament--"the Super Bowl of electronic sports", OpenAI challenged the top human player, Danylo Ishutin, to play against its AI.

The AI had been trained to play the real-time strategy game *Dota 2*.

During the match, the AI was beating Ishutin so badly that he begged it to "Please, stop bullying me."

> OpenAI first ever to defeat world's best players in competitive eSports. Vastly more complex than traditional board games like chess & Go.
> — Elon Musk (@elonmusk) August 12, 2017

Elon Musk, an investor in OpenAI, acknowledged the victory as a significant breakthrough.

In 2018, DeepMind achieved similar progress with AIs learning to play First Person Shooters. These games involve the tactics and strategies of individuals cooperating within a 3D environment. DeepMind's AI achieved superhuman levels of game play playing capture the flag in Quake III.

AlphaStar: The inside story

The top human players fall to another self-taught AI created by DeepMind.

In 2019, DeepMind created an AI called AlphaStar which learned to play the real-time strategy game *Star Craft II*. It defeated the top-ranked human player in 10 out of 10 games.

Beyond Games

AI has achieved incredible results in games. But how useful is it? An AI that plays Chess or Starcraft at superhuman levels won't cure cancer or save the world.

But general purpose learning algorithms are broadly useful. They can be used in many domains beyond mere games. Google's GNMT, for example, learned how to translate languages, which helps travelers navigate foreign lands.

One of the latest projects of DeepMind is an AI called *AlphaFold*. AlphaFold is learning how to solve the incredibly complex problem of protein folding — a problem which if solved has applications in developing medicines, gene therapies, and perhaps even cures for cancer.

In a recent paper, AlphaFold achieved higher accuracy in predicting protein folding than any prior method.

> AlphaZero's creative insights coupled with the encouraging results we see in other projects such as AlphaFold, give us confidence in our mission to create general purpose learning systems that will one day help us find novel solutions to some of the most important and complex scientific problems.
>
> *— DeepMind Team*

In 2016, DeepMind applied their algorithm to the problem of cooling their data centers. The result: The AI was able to cut their cooling bill by 40%.

Learning algorithms could also play a key role in achieving human and superhuman intelligence. For we don't need to figure out how to build the brain of an adult human — we only need to figure out how to build the brain of a baby. From there, the AI can learn everything else.

Movement abilities of AI

The defining difference between animals and stationary life forms is that we can navigate through an ever-changing environment. This requires continual acquisition of information to decide how to best get from point A to point B, while avoiding the mouths of hungry predators.

You could say, the reason we have brains is because we move. There is an animal that demonstrates this.

Once it settles down the sea squirt 'eats' its own brain. Image Credit: Nick Hobgood

The sea squirt in its larval stage can swim, but it has no digestive system. Before it can eat it must find a suitable spot on the sea floor to attach itself and begin a metamorphosis. When it emerges from this process, it has a mouth an digestive system, but its sense organs and most of it's brain are gone!

> Once attached, the juvenile adult no longer needs the sense organs, nerve chord or even its tail, so it reabsorbs them. The brain vesicle is transformed into a cerebral ganglion, which only helps the stationary adult to feed.
>
> *— John Bishop*

For AI to succeed in the physical world, it must also master movement in a dynamic environment.

Much progress has been made in this area. AI now achieves super-human performance in the control of ground-based and aerial vehicles. Further, it has mastered fine motor control over bodies.

Driving

Today's AI systems drive better than the average human. They benefit from millisecond reaction times, 360-degree vision, and never getting distracted, tired, or intoxicated.

Google started developing self-driving cars in 2009 with a 10-year goal of creating the world's best driver.

> Our #tenyearchallenge has been building the world's most experienced driver. Thanks to two visionary @Google characters for getting us started & to the @Waymo One riders in #Phoenix we're serving. HBD #Waymo ?
> pic.twitter.com/Ew4fdXjM7c
> — John Krafcik (@johnkrafcik) January 17, 2019

By 2014, Google's autonomous vehicles had logged over 700,000 miles and had done so without causing a single accident. Today, in certain places, you can order one to pick you up.

Inside The City Where Waymo Tests Self-Driving Vehicles

In Phoenix, Arizona, you can summon one of Google's autonomous Waymo vehicles to take you somewhere.

Human drivers only gain experience from a single perspective over a single lifetime.

AI systems aren't subject to this limit. Data from every autonomous vehicle can be combined to form a shared pool of experience from which lessons may be learned. This new experience then develops a better AI.

With a million AI vehicles on the road, it gains a million years of driving experience each and every year. It would know every road in the country, how to drive in every weather condition, and it would have the practice of avoiding dozens of accidents each day. No human driver could come close to being so experienced.

Flying

In addition to ground vehicles, Google is working on drones under Project Wing.

Google Drones Can Already Deliver You Coffee In Australia

Denizens of Canterbury, Australia can have hot coffee delivered in 3 minutes by autonomous drones.

The goal of Wing is drone-based delivery. It is currently being tested in Australia, where people can place food and beverage orders through an app. In 2019, Project Wing became the first drone delivery company to obtain an *Air operator's certificate* from the Federal Aviation Administration. This allows Wing to operate as a commercial airline in US airspace.

In military aviation, fighter pilots are the best of the best — their lives depend on it.

Gene Lee is a former fighter pilot with decades of experience. He has flown and commanded thousands of air intercept missions and was a Battle Manager and adversary tactics instructor for the U.S. Air Force.

Lee has fought AI opponents in flight simulators for decades. But in 2016, he met his match: an AI pilot called *ALPHA* shot down Gene Lee in every encounter. Fortunately for Lee, it was only a simulation.

ALPHA was in his words, "the most aggressive, responsive, dynamic and credible A.I. I've seen to date."

> I was surprised at how aware and reactive it was. It seemed to be aware of my intentions and [was] reacting instantly to my changes in flight and my missile deployment. It knew how to defeat the shot I was taking. It moved instantly between defensive and offensive actions as needed."
>
> *— Gene Lee, Retired United States Air Force Colonel*

Nimble Robots

AI can drive and fly vehicles, but could it ever pilot something like a human body, with its hundreds of muscles and joints and its constant need for balance? A company called *Boston Dynamics* has shown that it can.

They built a robot called *Atlas*, with a humanoid body. It is bipedal and moves like a person. It can navigate over difficult terrain and manage obstacles.

In a recent demo, Boston Dynamics showed Atlas to be capable of much more. It can run, jump, do back flips, cartwheels, handstands, and somersaults. It performs acrobatic feats that few humans can.

More Parkour Atlas

The Atlas robot by Boston Dynamics performs a floor routine in the style of a professional gymnast.

Recognition abilities of AI

The human brain is a pattern recognition machine. The most evolved part of the mammalian brain is its outer layer, known as the cortex. In humans, 30% of the cortex is devoted to processing visual information.

The vision system is responsible for taking flickers of light picked up by cells in our retinas and constructing the rich 3D world we know — one filled with objects, people, and meaning.

Though it feels automatic, billions of your neurons work behind the scenes to make it happen.

It is no small feat to build a machine that replicates the functions of our most complex sensory system. To make something that can determine the people, places, and things contained in patterns of light.

These patterns of light, grids of pixels, are ultimately just lists of numbers: the intensity of red light here, the intensity of green light there, etc. This is true not just for computers but also the brain. The brain sees only the counts of synaptic signals from the optic nerves--cords containing 1 million nerve fibers each.

Stanford Vision Lab and Princeton University created ImageNet — a collection of 14 million annotated images covering twenty-thousand categories. The data set is used to train systems for object recognition.

Since 2010, the ImageNet Large Scale Visual Recognition Challenge has let teams compete to see who can build the best object recognition system trained on their data set.

In the first year, results were abysmal.

The winning team had an error rate of 26% — five times greater than the human error rate.

Progress was slow, but steady. By 2015, just five years after competitions began, Microsoft Research built a system that exceeded the accuracy of humans. It had an error rate of just 4.94%.

As of 2020 the top system has an error rate of 1.3%.

Microsoft has an online demo of their image recognition technology. Google even lets you upload your own photos and have them be processed by their AI.

Google's object recognition applied to a picture. Note that the AI highlights the region in the picture where the elephants are.

It's notable that these AIs were not programmed by hand. Instead they learn by example.

The AI is not told what properties to look for to tell cats apart from dogs. Instead the AI is given many examples of cats and dogs and told which is which. The AI learns the distinguishing features of cats and dogs on its own.

Creative abilities of AI

A defining trait of humans is our creativity. We have the ability to come up with new ideas, and express originality. We apply our creativity to create art, music, and inventions.

> **Noun**.
> Creativity: the use of the imagination or original ideas, especially in the production of an artistic work.
>
> *— Oxford Dictionary*

Can a machine ever have a creative spirit or an imagination?

Original thought from a machine sounds like a contradiction in terms — machines are designed and programmed by humans. We tell them what to do. How then could a machine be *creative*?

Despite the seeming contradiction, machines have been made that challenge humans in creative domains. There are now machines adept at making art. Machines have conceived of patent-worthy inventions. Arguably, some machines have even shown that they have an imagination.

These endeavors have spawned a new field in AI called *computational creativity*.

Art

In 2017, a team from Berkeley's AI Research laboratory created an image translation system called *CycleGAN*. Once trained, it could turn pictures of horses into zebras, or scenes of winter to summer.

A picture of Yosemite taken in winter "reimagined" by the AI as a summertime scene. Note the added foliage and melting of snow. Image Credit: CycleGAN

The creators of CycleGAN took it a step further. They decided to train the AI in the styles of various impressionist painters. Once trained, they provided a photograph to the AI and asked it to generate a painting in the style of Claude Monet, Vincent Van Gogh, Paul Cézanne or the Japanese style of Ukiyo-e.

Paintings created by an AI trained in the styles of different impressionist painters. Image Credit: Taesung Park

Their result was incredible: any image provided to the AI was converted to a painting in the style of Monet or van Gogh.

But this AI has merely learned to copy the styles of other artists. It may take skill, but what creativity is involved? A truly creative AI would form its own style.

This is exactly the challenge that Ahmed Elgammal, director of the Art and Artificial Intelligence Lab at Rutgers, aimed to solve. He created a new type of AI he calls a *Creative Adversarial Network* or (CAN). The goal of a CAN is to create novelty: for example artwork in styles different from what it's seen before.

Accordingly, artwork produced by a CAN leans towards abstract pieces. Elgammal said, "I am surprised by the output every time I run it." Below is a sample of the AI's creations.

Some of the art pieces generated by Elgammal's Creative Adversarial Network. Image Credit: Ahmed Elgammal.

> We measured the difference in responses towards the human art and the machine art, and found that there is very little difference. Actually, some people are more inspired by the art that is done by machine.
>
> *— Ahmed Elgammal*

We need not wonder whether some day human artists will have to compete against machines. That day has already passed. In 2018, the auction house Christie's became the first to list artwork created by an AI.

The portrait fetched $432,500 — 45 times its estimated price!

“Portrait of Edmond De Belamy” was generated by an Artificial Intelligence.
Image Credit: Obvious art collective

Hugo Caselles-Dupré, the human who used AI to create the art, said, “We found that portraits provided the best way to illustrate our point, which is that algorithms are able to emulate creativity.”

Music

Musical expression is not uniquely human. Whales and birds also sing. But could an AI create music?

In 2017, a team from Microsoft created *BachBot* — an AI trained to compose music in the style of Johann Sebastian Bach. To test it, the researchers built a website, BachBot.com, where people can listen to samples and guess whether it was something created by Bach or by BachBot.

Out of 2,336 participants, the fraction of people that could tell BachBot from the real Bach was only 1% higher than could be expected from random guessing. In other words, people had a negligible ability to tell BachBot from the real thing. You can try it yourself at BachBot.com.

> As someone who makes music, the idea that code can do what I do is freaky. It’s unnerving to think that an algorithm can make a not-terrible song in minutes and that AI is getting in on creative turf we categorize as distinctly human.
>
> *— Dani Deahl, writer for The Verge*

AI is increasingly being used in the music industry, and has applications in writing lyrics, melodies, and even predicting sales that a musical piece will generate.

Break Free Official Music Video - Composed with AI | Lyrics by Taryn Southern

The musical melody behind “Break Free” was composed by an AI.

There is an online demo that allows you to use an AI to generate song lyrics with different topics, genres, and tones. When given the topic of **cat**, with a **pop** genre and **happy** mood, it produced:

> **Verse 1**
> You rise so that you embrace me
> Like a cat downs the cat with the nose up
> He got this cat, he got this cat
> He got this cat, he came home
> **Pre-Chorus**
> I wish I was a cat that can do all the things a cat can do
> I wish I was a cat that can leap from the ceiling
> **Chorus**
> So how did you manage to open up my heart, and still feel so alone
> The snow is falling, and my cat is hiding
> From a storm inside my heart from not sleeping
> And cats with silver chains

Recent evidence suggests that AI has even developed a *taste* for music. At least, it seems to understand how much humans will like a given piece of music.

The startup *Hyperlive* claims to have created an AI that predicts sales figures with 83% accuracy.

AI predicted sales figures in Yellow, and actual sales figures in White. Image Credit: Hyperlive.fm

Their system predicted sales figures for 200 tracks by 10 major artists. Their technology works by modeling the "neurological, physiological and behavioural responses" the music will invoke in the listener's brain.

Given its accuracy, we must wonder whether machines can have an aesthetic sense.

Imagination

Imagination is the ability to generate new ideas and imagery out of nothing.

In the last few years, the capacity for machines to imagine has not only arisen, but in a short time has vastly outstripped the abilities of human imagination, at least in terms of accuracy and realism.

Take for example, the following photo of a person. This person may look real, but in truth she does not exist. She was never born and never stopped to take a photo. There is no mind or soul behind her eyes.

A non-existent person imagined by an AI. Image Credit: thispersondoesnotexist.com

The photo and person were imagined by an AI. Created on-the-fly by visiting thispersondoesnotexist.com. You too can visit this site and refresh it to get an endless supply of photos of people who don't exist.

DeepMind extended this technology beyond faces. It has the ability to invent realistic pictures of cheeseburgers, dogs, butterflies, and landscapes.

There are also AIs that can be inspired to imagine. One takes rough sketches of cats, buildings, or shoes and turns them into photo-realistic interpretations. Another AI turns textual descriptions into images:

Photographs of Birds and Flowers generated from textual descriptions.
Image Credit: Generative Adversarial Text to Image Synthesis

This technology has improved beyond photographs into the realm of video.

It can be used to swap out actors in movies, such as replacing Arnold Schwarzenegger with Sylvester Stallone in Terminator 2, or used in the popular trend of putting Nicolas Cage in every movie.

With the latest technology, an AI can take a single photograph and bring it to life by predicting the future. A team from MIT's AI Lab trained an AI to predict the next few seconds of video from a single picture.

In 2019, Samsung's AI Lab created *Few-Shot*. It can take a single photo, and imagine it as a living, emoting, and talking person. To demonstrate the power of needing only a single photo they used the Mona Lisa.

AI brings Mona Lisa to life

The Mona Lisa is brought to life through the imaginative power of the AI called Few-Shot.

Invention

Invention is another expression of human creativity. Inventing new tools and technologies we have transformed our way of life. Human creativity brought us from tribes people to having every convenience of modern society.

But humans are no longer the sole possessors of ingenuity.

The computer scientist John Koza created the field of *genetic programming*.

Genetic programming is inspired by the techniques of biological evolution. It enables a program to update and adapt itself in the search for more optimum solutions to problems.

Koza used these techniques to build the *Invention Machine*.

The invention machine has optimized and rediscovered the designs of antennae, circuit boards, lenses and factories. It was able to, in Koza's words "automatically synthesize complete designs for six optical lens systems that duplicated the functionality of previously patented lens systems."

In 2005, one of the designs created by the Invention Machine was granted Patent No. 6,847,851 by the U.S. Patent and Trademark Office.

The rise of AI inventors is creating problems for patent offices around the world.

In 2018, the European Patent Office rejected two inventions where an artificial intelligence was the sole inventor on the grounds that "Machines do not have legal personality and cannot own property."

> There are machines right now that are doing far more on their own than to help an engineer or a scientist or an inventor do their jobs. We will get to a point where a court or legislature will say the human being is so disengaged, so many levels removed, that the actual human did not contribute to the inventive concept.
>
> *— Andrei Iancu, director of the U.S. Patent and Trademark Office (2020)*

We've already seen that AI can invent new strategies and tactics in games like Chess and Go.

> We all expect machines to play very solid and slow games but AlphaZero just does the opposite. It is surprising to see a machine playing so aggressively, and it also shows a lot of creativity.
>
> *— Gary Kasparov, Grandmaster and former World Chess Champion*

> I thought AlphaGo was based on probability calculation and it was merely a machine. But when I saw this move I changed my mind. Surely AlphaGo is creative.
>
> *— Lee Sedol, 9-dan professional Go player and former champion, referring to 'move 37'*

Reasoning abilities of AI

In 2004, Ken Jennings made game show history by winning 74 consecutive games on the quiz show *Jeopardy!*.

No one doubts the superior memory of today's computers. They far outpace the brain in terms of speed and accuracy. But storing facts, and knowing how to apply them is a whole other problem.

That's why IBM saw it as a great challenge to make an AI that could beat world champions at Jeopardy! — a game whose clues often require solving complex verbal puzzles.

In 2009, Jennings received a call from Jeopardy's producers. They asked, "IBM tells us they want to build a supercomputer to beat you at 'Jeopardy.' Are you up for this?"

In his day job, Jennings was a computer programmer. He knew computers lagged far behind humans at these kinds of problems. He thought, "this is going to be child's play. Yes, I will come destroy the computer and defend my species."

> People don't realize how tough it is to write that kind of program that can read a Jeopardy clue in a natural language like English and understand all the double meanings, the puns, the red herrings, unpack the meaning of the clue.
>
> *— Ken Jennings*

IBM's Jeopardy! playing AI *Watson* handily beats its world-class human opponents.

But when put to the test in 2011, neither Jennings, nor his compatriot on the side of humanity, Brad Rutter, could defeat Watson. The AI ended the three-game match with a score of $77,147 — more than the combined scores of its human opponents.

Accepting defeat, Jennings wrote under his final answer “I, for one, welcome our new computer overlords.”

Humans often judge the reasoning abilities of other humans through standardized multiple-choice tests.

Researchers from the Allen Institute for Artificial Intelligence published results on an AI test taker called *Aristo*. In just three years time, Aristo went from flunking science tests to acing them.

In 2016, Aristo could only answer 59.3% of the questions on the 8th grade N.Y. Regents Science Exam correctly but by 2019 Aristo advanced to answer over 90% of questions correctly. This test had questions like:

> **Which equipment will best separate a mixture of iron filings and black pepper?**
> (1) magnet (2) filter paper (3) triple-beam balance (4) voltmeter
> **Which form of energy is produced when a rubber band vibrates?**
> (1) chemical (2) light (3) electrical (4) sound
> **Because copper is a metal, it is**
> (1) liquid at room temperature (2) nonreactive with other substances (3) a poor conductor of electricity (4) a good conductor of heat
> **Which process in an apple tree primarily results from cell division?**
> (1) growth (2) photosynthesis (3) gas exchange (4) waste removal

According to the project leader Peter Clark, “Even five years ago, computers had a lot of difficulty understanding what was written in text. Thanks to a rapid progression of advances, we now have AI systems that are much better able to understand language.”

In the next section, we will look at the jobs that AI is about to take.

Artificial Intelligence and Jobs

AI that performs at the human-level in all areas will have profound societal, social, and economic impacts. More limited AI has already had far-reaching economic effects.

Horses made their living through their strength. They couldn’t find new jobs once stronger machines like automobiles and tractors entered the scene. The horse population peaked in the early 1900s.

What will happen to humans when a more creative, more intelligent, thinking machine enters the scene?

An AI can operate 24/7. It can work faster, more accurately, and for less money. The AI worker requires no training time. It never quits, takes leave, or gets sick. What chance do we have?

AI could be the gateway to a new golden age. One where we free ourselves from tedious labor and pursue our dreams. On the other hand, our economic obsolescence could spell our doom.

Let's review the jobs already lost to automation, and consider which jobs are under an imminent threat of replacement by AI.

Jobs Taken by AI and Automation

Many jobs have been automated away, or substantially reduced by AI, software, or machines:

- **Bank Tellers** – ATMs
- **Toll Collectors** – Toll machines, Toll transponders (E-ZPass, I-PASS)
- **Editors** – Spelling and Grammar check software
- **Personal Assistants** – Alexa, Siri, Google Home
- **Supermarket Cashiers** – Self-checkout machines, Amazon Go Stores
- **Warehouse Workers** – Robots
- **Assembly Line Workers** – Robots
- **Travel Agents** – Booking Websites
- **Tax Preparers** – Tax Software

Human-level general intelligence threatens all jobs. Someday soon, we may all feel as these workers felt.

> I remember thinking, this is it. I felt obsolete. I felt like a Detroit factory worker of the '80s seeing a robot that could now do his job on the assembly line. I felt like quiz show contestant was now the first job that had become obsolete under this new regime of thinking computers.
>
> *— Ken Jennings*

Jobs Threatened by AI

A 2019 report estimates 36 million jobs are threatened by emerging AI technologies.

When we think of jobs ripe for replacement, we usually think of low-skilled, jobs that pay little and need little training. While this has been true in the past, new systems threaten highly-paid and highly-skilled jobs.

Grocery store cashiers and truck drivers are not the only ones that need worry about robots taking their jobs. Jobs threatened by AI include: lawyers, radiologists, anesthesiologists, pharmacists, scientists, airline pilots, hedge fund managers, and Hollywood actors.

AI in Law

The field of law includes paralegals, lawyers, and judges. All these jobs face the threat of automation.

Paralegal – *average salary: $51,740*

Paralegals work as assistants and interns at law firms. Many go on to become lawyers themselves. They help lawyers prepare cases and search for relevant facts in documents and case law.

A technology called *electronic discovery* now automates much of this work. It is faster, cheaper, and more thorough than manually sifting through the millions of documents that might be involved in a case.

> From a legal staffing viewpoint, it means that a lot of people who used to be allocated to conduct document review are no longer able to be billed out. People get bored, people get headaches. Computers don't.
>
> *— Bill Herr, a lawyer who previously filled auditoriums with lawyers to pour over documents*

Lawyer – *average salary: $94,495*

Lawyers average seven years of education after high school: four for an undergraduate degree, plus three more at law school. Despite this, even lawyers face the threat of automation.

In 2018, twenty top corporate lawyers were pitted against an AI Lawyer called *LawGeex*.

The lawyers and the AI were asked to perform a standard lawyer function: reviewing a contract to identify potential issues. Human lawyers achieved an accuracy of 85% and it took them an average of 92 minutes to complete the task. The AI system LawGeex achieved an accuracy of 94% and finished in 26 seconds.

One of the lawyers in the competition said, "We are seeing disruption across multiple industries by increasingly sophisticated uses of Artificial Intelligence. The field of law is no exception."

Judge – *average salary: $120,090*

We rely on judges for both their knowledge of the law and their fairness.

The country of Estonia plans to use an AI to settle small claims cases. Both parties will upload their documents to be processed by an AI to give a decision. The AI's decision can be appealed to a human judge.

In the US, some states are using algorithms to decide prison sentences. It remains to be seen whether AI will result in fairer verdicts and sentences, but unlike human judges, AIs will only get better over time.

AI in Medicine

The field of medicine includes some of the most-highly paid professions. It includes pharmacists, radiologists, anesthesiologists, and surgeons. All these jobs are threatened with replacement by intelligent machines.

Pharmacist – *average salary: $128,090*

Pharmacists are responsible for filling prescriptions. Robotic pharmacists, such as those developed by SwissLog and the University of California San Francisco, fully automate identification, counting, and dispensing of pills.

The *PillPick* system fully automates a pharmacy. Image Credit: SwissLog

> A nurse made an error of putting the decimal point in the wrong place and we overdosed a patient, and at that point, we made a commitment to that we didn't ever want that to happen again.
>
> *— Mark Laret, CEO of the University of California San Francisco Medical Center*

In testing at the University of California San Francisco, the automated pharmacist filled 350,000 prescriptions with 0 errors. Human pharmacists have an error rate of 5 errors per 100,000 prescriptions.

Surgeon – *average salary: $255,110*

Surgeons not only require medical expertise, but adept hands and keen eyes.

New robotic tools are deployed in thousands of operating rooms across the world. They enhance the capabilities of human surgeons. Robotic surgeons like da Vinci Surgical System give surgeons 4 arms, magnified 3D vision, hand stabilization, fine motor control, and 360-degree articulating wrists.

da Vinci Surgical System: Surgery on a grape

The incredible precision of robot assisted surgeries enables otherwise impossible operations. For example: heart surgeries that don't require cracking open the patient's ribs.

While today's robotic surgeries are performed with a human surgeon at the helm, the company *Digital Surgery*, now part of Medtronic, aims to add enhance robotic surgeons by connecting them with AI.

> We're not anywhere near playing grand-master chess. But the computers are at the level of a medical-school student.
>
> *— Jean Nehme, surgeon and founder of Digital Surgery*

Anesthesiologist – *average salary: $261,730*

Among medical specialties, anesthesiologists are among the best paid. The job requires expert skill and constant diligence. They must give patients just enough sedatives to render them insensitive but not so much that they stop breathing.

Johnson & Johnson built a machine to automate patient sedation called *Sedasys*. It was among the first devices to automatically administer sedatives while monitoring the patient's condition.

A Washington Post article about Sedasys "provoked an outpouring of messages from anesthesiologists and nurse anesthetist who claimed a machine could never replicate a human's care or diligence."

But others were less certain. One patient asked a friend and anesthesiologist what he thought about it.

His response: "That's going to replace me."

Since the Sedasys, more advanced systems have been built. One example is the *iControl-RP*. It's a *closed-loop* system — it receives no external instructions. It makes all dosage decisions on its own. The iControl-RP not only monitors the patient's vital signs, but it even monitors the patient's brain waves.

A recent study found that patients sedated by closed-loop systems recovered faster.

> I have no doubt that closed-loop (i.e. robotic) anesthesia is at least as good as the best human anesthesia. And that, for me, would be good enough to use it every day.
>
> *— Thomas M. Hemmerling, M.D., D.E.A.A.*

Radiologist – *average salary: $414,890*

Radiologists pour over patient scans to look for signs of disease and to diagnose patients.

In 2019, researchers compared the performance between an AI system and 101 professional radiologists.

In particular, they compared the accuracy and sensitivity of radiologists in screening mammograms for signs of breast cancer. In processing 28,000 mammogram images, researchers found the AI performed better than a majority of the human radiologists.

> Some medical students have reportedly decided not to specialize in radiology because they fear the job will cease to exist.
>
> *— Harvard Business Review*

AI in Science

The field of science attracts some of humanity's greatest minds. Their discoveries drive progress.

But could machines ever be so bright as to make new scientific discoveries?

Medical Scientist – *average salary: $88,790*

Drug discovery can be lucrative, but it is also costly. Out of 5,000 drugs that show promise in preclinical trials, only 5 make it to human testing. Of those 5 that make it to human testing, only 1 makes it to market.

The process of going through laboratory testing, animal testing, phase I, II, and III human clinical trials, and finally obtaining FDA approval takes an average of 12 years and $2.6 billion.

Arcinova's drug research labs in Alnwick, England. Image Credit: Arcinova

It's little wonder why companies might invest in technologies to streamline the process.

This led to the creation of an AI called *Centaur Chemist* — a system to automate the process of drug design.

In 2020, human clinical trials began on the compound *DSP-1181*. This drug was identified by Centaur Chemist for the treatment of obsessive-compulsive disorder. The AI discovered it in 12 months--a quarter of the time normally required by human drug research teams.

Physicist – *average salary: $122,220*

Physicists use observations to develop models that describe the world. In 2020, two MIT physicists, Silviu-Marian Udrescu and Max Tegmark built an AI that could do the same — it rediscovered known laws of physics.

> We just posted a new AI paper on how to automatically discover laws of physics from raw video with machine learning. For example, we feed in the video below of a rocket moving in circles in a magnetic field, seen through a distorting lens, and our code automatically discovers the Lorentz Force Law.
>
> *— Max Tegmark*

AI in Finance

Few fields are more lucrative than those of finance, which includes traders and hedge fund managers.

Securities Trader – *average salary: $98,770*

Since a 2001 paper by IBM showed algorithms "consistently obtain significantly larger gains from trade than their human counterparts," the fraction of trades performed by humans has dwindled.

The fraction of trades performed by algorithms as a percentage of the total. Image Credit: Wikipedia/Tertius51

As of 2012, for every 7 trades made, 6 were made by computers. Today's trading floors now sit mostly empty. All the real activity takes place in data centers and server rooms.

Fund Manager – *average salary: $129,890*

It's a great responsibility to be a mutual fund or hedge fund manager. They are entrusted with billions of dollars of assets. But recent AIs are giving fund managers a run for their money.

In the past few years, hedge funds run by AI have outperformed the returns of human-run funds. Image Credit: Preqin Pro

AI in Entertainment

The field of entertainment employs a host creative types and supporting personnel: actors, stuntmen, directors, costume teams, makeup artists, screen writers, film crews, as well as set and prop designers.

All of these categories of employment could be rendered obsolete by advances in AI.

We have already seen AI that can take textual descriptions and turn them into images. We've also seen AIs that can turn images into video and realistic body motions. AI can already swap out likenesses of actors, or change background scenes (e.g. from winter to summer or day to night).

Might we see a point where the tools of AI create feature-length films digitally? If so there would be no need of actors, costume teams, film crews, traveling to destinations, obtaining permits, nor waiting for the right weather or time of day. It gives the creator unlimited control — constrained only by imagination, not budgets.

But it would leave many in Hollywood without a job.

Digitally-rendered movies are not new. They have been around for decades, the first being Toy Story in 1995.

Toy Story (1995) was the first full-length movie generated entirely on a computer. Due to the limited power of the computers and algorithms of the time its graphics lacked the realism expected by today's standards. Image Credit: Pixar

Only recently have digitally-rendered movies begun to approach photo-realism.

Note the level of detail in the wood grain of the floor in Toy Story 4 (2019). Image Credit: Pixar

We've also seen the rise of entirely digital actors, like *Gollum* from The Lord of the Rings (2001).

LOTR - Best of Gollum

Recently, Hollywood used digital recreations of deceased actors. The actor Peter Cushing, who died in 1994, was digitally resurrected to star as *Grand Moff Tarkin* in Rogue One (2016).

AI and algorithms will allow, at the click of a button, the appearance of a digital actor to be tweaked, and all video scenes with that actor will be automatically updated — without having to do any reshoots. Aided by AI, creating movies could one day be a hobby accessible to anyone.

AI in Transportation

The field of transportation employs millions. It includes those who drive or fly trucks, delivery vans, buses, trains, planes, taxi cab drivers, and ride-share vehicles.

Google's Waymo, G.M. Cruise, Tesla, Apple Car, Baidu, and some 35 other companies are working to develop autopilot technology that will one day put all these transportation workers out of work.

Truck Driver – *average salary: $42,260*

Truck Driver is one of the most common professions in many states.

Not counting managers or salespersons, 'Truck Driver' is the most common job in most states. Image Credit: NPR

> This week, we'll start driving our Chrysler Pacificas and long-haul trucks in Texas and New Mexico. These are interesting and promising commercial routes, and we'll be using our vehicles to explore how the Waymo Driver might be able to create new transportation solutions. pic.twitter.com/uDqKDrGR9b
> — Waymo (@Waymo) January 23, 2020

In the not too distant future, the job of long-haul Truck Driving could become a thing of the past.

Airline Pilot – *average salary: $161,280*

Autopilot systems exist on every large commercial jetliner. They not only fly but can land the plane.

The standard procedure for most airlines is the use of automation for much of the flight.

> "Let the computer do it because it can do a better job than a person."
>
> *— Paul Robinson's account of the general advice given to pilots*

But today's autopilot systems don't taxi the plane, nor do they take off. Even though autopilots can land, the majority of landings are performed manually. New autopilot systems aim to change that.

ATTOL: Autonomous Taxiing, Take-Off and Landing test flight

Airbus is currently testing fully autonomous taxiing, take off and landing (ATTOL) technology.

Boeing is also working on a fully autonomous airplane.

AI in Writing

Journalist – *average salary: $62,400*

We've seen that GPT-2 can write in almost human style.

In 2019, John Seabrook, writer for *The New Yorker*, traveled to meet with the people at OpenAI behind GPT-2. He wanted to answer the question: Can a Machine Learn to Write for the New Yorker? He concluded:

> The brain is estimated to contain a hundred billion neurons, with trillions of connections between them. The neural net that the full version of GPT-2 runs on has about one and a half billion connections, or "parameters." At the current rate at which compute is growing, neural nets could equal the brain's raw processing capacity in five years. To help OpenAI get there first, Microsoft announced in July that it was investing a billion dollars in the company, as part of an "exclusive computing partnership."
>
> *— John Seabrook, staff writer for The New Yorker*

Microsoft may be well on their way to getting there first. In 2020, they announced that they are firing nearly 80 people and replacing them with AI that will automate the curation of their news pages.

Screen Writer – *average salary: $73,090*

AI could one day entirely automate the process of bringing a movie script to life, But someone still has to come up with the stories. Right?

Perhaps not. In 2016, an AI wrote a short screenplay that was turned into a film called *Sunspring*:

Sunspring | A Sci-Fi Short Film Starring Thomas Middleditch

Sunspring: A science fiction film written by an AI.

Though the dialogue is laughably bad, this represents the floor. AI storytellers will only improve from here.

What Jobs are Safe?

In the long run, the rise of sophisticated AI and robotics means no job is completely safe.

However, in the short run, some jobs seem safer than others. These include jobs that need a human touch, or a human personality to provide inspiration, comfort, compassion, or persuasion. These jobs include:

- Childcare workers, teachers, and nurses
- Waitstaff and bartenders
- Social workers and therapists
- Airline pilots (to reassure passengers who wouldn't fly without human pilots)
- Managers, sales people, CEOs, and public speakers

There are also some that by convention or law require humans to fill them, such as:

- Politicians and other holders of public office
- The Clergy
- Professional athletes
- Live performers (in plays, music, comedy)

Finally, there are some jobs that might survive until AI attains a human-level understanding of the world:

- Comedians (to know what people find funny requires knowing how people think)
- Movie critic
- Writing content that is long and coherent (screen plays, novels, books)
- AI Programmers (This job will exist until AI programmers make AI smart enough to replace themselves)

The list of things humans can do better than AI is short and shrinking. In this final section, we will examine the near and far future capabilities of artificial intelligence.

The Future of Artificial Intelligence

What would we get if we put all the previous AI's in the same computer?

We would get an AI that could hold a conversation in any of 100 languages over the phone. It would beat everyone at Jeopardy!, Chess, Go, Poker, and in many video games. The AI would be able to recognize any object or face — even detect cancer better than most doctors. It would also be accomplished and creative, having invented things, discovered laws of physics, and identified new drugs. The AI could compose as well as Bach and paint as well as van Gogh. It would also be seen as highly original in its own artistic styles.

The AI would be a modern day Renaissance machine — with a collection of skills unrivaled by any human.

> 'Let's put all these together,' and then it will be smart.
>
> *— Marvin Minsky*

At one point in history, it was possible for a single person to be well versed in all of human knowledge. But as the domains of science specialized and human knowledge grew, this hasn't been possible for centuries.

However, an AI could hold all human knowledge in its head. Every encyclopedia, every book and science article ever published, even the contents of every web page on the Internet. With the ability to cross-reference, analyze, find patterns in, and summarize written text, as GPT-2 does, it is only a matter of time before we build an entity with an understanding of the collective knowledge of humanity.

Algorithms of Intelligence

In 2000, Marcus Hutter, now at DeepMind, discovered an algorithm for *universal artificial intelligence*, called AIXI.

Universal Artificial Intelligence can be described by a single equation. Image Credit: Marcus Hutter

When followed, this algorithm achieves perfect and optimal intelligence. It works by computing every possibility resulting from every available course of action, then choosing the action that brings it closest to its goals.

If the goal is to win a chess game, AIXI computes every possible future chess game resulting from each of the possible moves, and chooses the move with the most wins and fewest losses among its future possible games.

The algorithm has one downside: It's uncomputable — it needs unlimited computational resources to solve.

Hutter's formula does offer an important lesson: *intelligence isn't complicated* — what's hard is getting intelligence with limited computing power. Achieving intelligence with limited resources requires finding and taking shortcuts--matching patterns and applying heuristics--in place of brute compute.

It seems we now know how to do this. DeepMind's AlphaZero is an algorithm for learning. OpenAI's GPT-2 is an algorithm for understanding. BAIR's CycleGAN is an algorithm for creativity.

With the algorithms of intelligence, only one thing stands in the way of human-level AI: raw computing power.

Neural Networks

AlphaZero, GPT-2, CycleGAN, Microsoft's speech and object recognition, Google's GMNT translation, *all of them*, are based on the technology of *artificial neural networks* — which were themselves inspired by the brain.

Neural networks have been used since the 1960s, and were a primary topic of interest by AI researchers in the 1980s. But the field of AI stagnated when neural networks failed to do anything useful.

Their failure, however, was not due to having the wrong principles, but not having enough computing power.

Trying to build useful AI with computers of the 1980s was like trying to get to the moon with a bottle rocket. The principle is right, but the power was not there.

With increasing computing power, larger neural networks, with many more neurons and many more layers became computationally feasible. These networks are known as *deep neural networks*.

In an artificial neural network, inputs, such as pixels of an image, are processed by successive layers of neurons until a final output is reached. Deep neural networks simply use more layers. Image Credit: RBC

Each layer of the network can separately identify unique features. This was revealed in 2015 by Google's DeepDream, a visualization of what each layer in a deep neural network sees.

Deep Dream: The Grove - Exploring Layers

What each layer "sees", each extracts different elements of the picture. *Computers can hallucinate*.

Deep neural networks of today can have dozens of layers and millions of neurons. Recent breakthroughs aren't a result of discovering some new paradigm of intelligence but the result of building faster processors and larger data sets. The current AI revolution is almost entirely a result of increased computing power.

Thanks to pioneers like Marvin Minsky in the 1980s and Jürgen Schmidhuber in the 1990s we've had the right methods for decades. However, only recently have we had the computing power to run and train networks large enough to manifest interesting and intelligent behaviors.

Now, as the power of our best computers approaches the raw computing power of the human brain, we're seeing increasingly human-like achievements in AI.

How much time is left before our computers, and the AI they power, eclipse us?

Computing Trends

It is the nature of technology to improve, and its speed of improvement has been *exponential*.

Exponential growth occurs whenever the amount of growth is proportional to something's magnitude. For example, the interest paid to a bank account is proportional to its balance.

Whenever there is exponential growth there will be a constant doubling time.

In 1965, Gordon Moore, who would later found Intel, noticed that the power of computers was doubling almost every year. The trend could be traced back well before the sixties, and it has continued ever since.

Since 1900 we've had a 10^{18}-fold increase in the performance of computers. Image Credit: Steve Jurvetson

Exponential growth occurs whenever there is a feedback mechanism. For example, the better our knowledge about building computers, the faster and more powerful computers we can make.

The faster computers we make, the faster we can gather and process information to expand our knowledge, which includes new knowledge about building computers. The cycle repeats, and feeds on itself.

We observe exponential growth in every measure of our knowledge — in the number of new patents filed, in the number of scientific articles published, and in the total amount of digital data stored.

The number of patents (a proxy for technological progress) is exploding exponentially. Image Credit: Our World in Data

Today, is just a starting point. If the power of computing technology continues to double every year, then in a decade our computing power will expand by 1,000 times: 1,000 times faster and 1,000 times more memory.

Applied to AI that means each decade AI gets *1,000 times smarter*.

> There will be about thirty doublings in the next 25 years. That's a factor of a billion in the capacity and price performance over today's technology, which is already quite formidable.
>
> *— Ray Kurzweil*

Should AI ever catch up to us, it won't stay at our level for long. It will go soaring past us.

An Intelligence Explosion

We feel as though we are in the midst of something big — a paradigm shift, the next stage of life, something.

Recordings of these feelings go back to at least the late 1950s.

> One conversation centered on the ever accelerating progress of technology and changes in the mode of human life, which gives the appearance of approaching some essential singularity in the history of the race beyond which human affairs, as we know them, could not continue.
>
> *— Stanislaw Ulam, in 1958, recounting a conversation with John von Neumann*

Danny Hillis, who worked with Minsky's at the MIT AI Lab compared it to being in the middle of an S-curve.

Danny Hillis: Back to the future (of 1994)

> So the first steps of the story that I told you about took a billion years a piece. And the next steps, like nervous systems and brains, took a few hundred million years. Then the next steps, like language and so on, took less than a million years. And these next steps, like electronics, seem to be taking only a few decades.
>
> The process is feeding on itself and becoming, I guess, autocatalytic is the word for it – when something reinforces its rate of change. The more it changes, the faster it changes. And I think that that's what we're seeing here in this explosion of curve. We're seeing this process feeding back on itself.
>
> *— Danny Hillis*

We are alive at a most exciting time in history. But does it pose a threat to human life as we know it?

The Doomsday Equation

What might this next stage of life be and what possibilities will it bring? How much time do we have left before it enters the world stage? Various signs, from different fields, all suggest it could be as near as a few decades.

Heinz von Foerster was versed in the fields of computer science, neurophysiology, mathematics and philosophy. The Pentagon funded von Foerster to create and lead the *Biological Computer Laboratory*, where he pioneered the field of cybernetics.

He published over 200 papers in his career, but his most famous is his 1960 *Doomsday Equation*.

Von Foerster's doomsday equation was the result of analyzing human population growth trends. He and his students gathered and analyzed the human population size over the previous two thousand years. They discovered it is growing faster than an exponential rate.

The growth was not exponential, but *hyperbolic*.

While an exponential trend doubles at a constant rate, von Foerster found the time between doublings was shrinking. They plotted when this doubling time was projected to reach zero — a time where human population would, if it followed this trend, shoot to infinity. They arrived at the following projection: **2027 A.D. ± 5.5 years**.

A similar pattern was discovered in economics. The economic historian James Bradford DeLong collected data to estimate world GDP over the previous one million years. Again, when plotted, it showed a trend of a decreasing time between successive doublings.

It suggested a point in the early 21st century when the doubling time of the economy would reach zero.

The productive capacity of the world economy is shooting towards infinity. Image Credit: Our World in Data

Two researchers created a model of population, technology, and inventors to estimate world technological development over time. They concluded:

> Extremely simple mathematical models are shown to be able to account for 99.2–99.91 per cent of all the variation in economic and demographic macrodynamics of the world for almost two millennia of its history.
>
> *— Andrey Korotayev and Artemy Malkov*

The trend of the data is so clear and consistent that someone in ancient Rome or the middle ages with the data of their time could have predicted this trend would reach its end sometime during the 21st century.

A Singularity in History

> An analysis of the history of technology shows that technological change is exponential, contrary to the common-sense 'intuitive linear' view. So we won't experience 100 years of progress in the 21st century — it will be more like 20,000 years of progress (at today's rate).
>
> *— Ray Kurzweil, in 2000*

Trends in the growth of population, the economy, and technology all point towards an emerging Technological Singularity in the near future. This is a point when machine intelligence vastly outstrips human intelligence.

Once that occurs, humans will no longer be in the driver seat of technological development. Unconstrained by the human population of scientists, inventors, and technologists, the only limit on the speed of technological progress will be the computing resources available for AI-based scientists, inventors and technologists.

How much computing power $1,000 can buy over time. Image Credit: Ray Kurzweil

As of 2018, our fastest supercomputer, Summit, exceeded the computational power of one human brain.

In a few decades of continued technological progress, our personal computers and smart phones will catch up to the computing power of Summit. Around this time, the total computing capacity of our machines will exceed the total computing power of all human brains.

> Essential historic developments match a binary scale marking exponentially declining temporal intervals, each half the size of the previous one [...] apparently converging to zero within the next few decades.
> The remaining series of faster and faster additional revolutions should converge in an Omega point expected between 2030 and 2040, when individual machines will already approach the raw computing power of all human brains combined. Many of the present readers of this article should still be alive then.
>
> *— Jürgen Schmidhuber, AI pioneer*

Superpowers of Superintelligence

Irving John Good was a mathematician who worked alongside Alan Turing using computers to break German codes in WW2. Good was one of the first to realize the implications of a machine that could improve itself.

> Let an ultraintelligent machine be defined as a machine that can far surpass all the intellectual activities of any man however clever. Since the design of machines is one of these intellectual activities, an ultraintelligent machine could design even better machines; there would then unquestionably be an 'intelligence explosion,' and the intelligence of man would be left far behind. Thus the first ultraintelligent machine is the last invention that man need ever make.
>
> *— Irving John Good, in 1965*

Such an intelligence would possess many attributes we might call *superpowers*. Nick Bostrom's book *Superintelligence* outlines six superpowers that superintelligent AIs might possess:

Superpower	Skills	Relevance
1. Intelligence amplification	AI programming, cognitive enhancement research, social epistemology development, etc.	• System can bootstrap its intelligence
2. Strategizing	Strategic planning, forecasting, prioritizing, and analysis for optimizing chances of achieving distant goal	• Achieve distant goals • Overcome intelligent opposition
3. Social manipulation	Social and psychological modeling, manipulation, rhetoric persuasion	• Leverage external resources by recruiting human support • Enable a "boxed" AI to persuade its gatekeepers to let it out • Persuade states and organizations to adopt some course of action
4. Hacking	Finding and exploiting security flaws in computer systems	• AI can expropriate computational resources over the Internet • A boxed AI may exploit security holes to escape cybernetic confinement • Steal financial resources • Hijack infrastructure, military robots, etc.

5. Technology research	Design and modeling of advanced technologies (e.g. biotechnology, nanotechnology) and development paths	• Creation of powerful military force • Creation of surveillance system • Automated space colonization
6. Economic productivity	Various skills enabling economically productive intellectual work	• Generate wealth which can be used to buy influence, services, resources (including hardware), etc.

Given its superpowers, a superintelligence aligned against humanity would be a curse. We would have little chance at prevailing against it.

However, a superintelligence on our side would be a blessing. It could cure any disease, design any technology, fix any problem, even end world hunger and poverty.

A superintelligence is how we could see 20,000 years of progress over the next 100 years.

> Everything that civilisation has to offer is a product of human intelligence; we cannot predict what we might achieve when this intelligence is magnified by the tools that AI may provide, but the eradication of war, disease, and poverty would be high on anyone's list. Success in creating AI would be the biggest event in human history. Unfortunately, it might also be the last.
>
> *— Stephen Hawking*

In two decades, we've seen the rise of intelligent machines. Machines that are creative, that learn, that fool us into thinking they are human. If so much progress can come in such a short period, what will be possible over the coming centuries as computing technologies continue to grow exponentially in power?

Limits of Intelligence

There are some limits even superintelligence cannot overcome.

For example, limits like the speed of light and matter densities of black holes. Physical laws imply physical limits on the processing speed, data density, and energy efficiency of computers.

(See: "How good can technology get?")

However, the ultimate physical limits on computation are extraordinary. *Bremermann's limit* bounds the speed of the fastest possible computer to 10^{50} operations per second per kilogram of mass.

The Summit supercomputer achieves about 10^{16} operations per second per kilogram, meaning we are about 10^{34} off from building the best possible computers. Computing technology could double another 112 times before we would reach this limit.

The best physically possible computing material is referred to as computronium. It is the stuff of science fiction, but we can use the known physical bounds to speculate about what computronium could do.

Humans often consider themselves as the pinnacle of intelligence, but in truth, even the combined intelligence of all human brains put together is but a speck next to what is possible:

A graph of the orders of magnitude that separate possible intelligences. Image Details: Google Docs

On the above chart, each increase by 1 represents a 10-fold increase in computing power.

The average pocket calculator can perform 10 operations per second and so is at a '1' on the chart. A modern smart phone that performs one trillion (10^{12}) operations per second, sits at '12'.

Chimps, Humans, and the Summit super computer are all around the exaop (10^{18}) scale, so all three sit around '18'. The combined power of all computers in the world is around 10^{21} operations per second, and all 7 billion human brains together is at (10^{28}).

On this scale, all of humanity put together sits halfway between a pocket calculator, and a *Matrioshka brain* — a hypothetical computer powered by a star, which achieves 10^{48} operations per second when operating at Laundauer's limit of computing efficiency.

Though much smaller, 1 kilogram of computronium, operating at physical limits, has 100 times the power of this star-powered computer, achieving 10^{50} operations per second.

Jupiter weighs about 10^{27} kilograms. Converting the entire mass of Jupiter into computronium would yield a computer capable of 10^{77} operations per second.

Ultimately AI could possess so much computing power that it could explore the entire evolutionary history of other worlds and civilizations in a fraction of a second through computer simulation.

Consider that in the history of humanity around 100 billion humans have lived. If each person lived an average of 40 years this amounts to a total of 4 trillion years of human experience.

Four trillion years is roughly 10^{20} seconds. Since each second of human brain activity involves 10^{18} operations, then all experiences had by all of humanity can be produced with 10^{38} operations.

An AI with 1 kilogram of computronium could experience all of humanity in a trillionth of a second.

A Jupiter brain of computronium could in each second, dream the dreams of 10^{39} civilizations.

Is our own existence within such a dream?

The potential power of future computers raises many questions, such as "Are we living in a computer simulation?", "Is it possible to create new universes?", and "Is life insignificant in the grand scheme of things?"

Conclusions

We began this article with the question of *When will AI take over?*

Are we any closer to an answer?

Are Intelligent Machines Possible?

In the recent past, one could make the argument that machines would never be intelligent, "Machines can only do what humans programmed them to do."

This argument no longer holds water.

The modern approach to AI relies on machines that teach themselves. Since they aren't given explicit instructions, machines have learned new strategies and styles no human programmer conceived.

One of the earliest tests of machine intelligence was proposed by the father of computing, Alan Turing. In his 1950 paper "Computing Machinery and Intelligence," Turing proposed *The Imitation Game*, now more commonly known as the *Turing test*.

Turing suggested we would know machines are intelligent when people can no longer reliably tell whether they are talking with a computer or with another human. By the measure of this test, and the results of Duplex, we have reached this point.

Machines are intelligent. They have achieved or are approaching human performance in nearly all tasks.

What does our Future Hold?

What role might humanity play in all of this? We certainly will play a key role in ushering in the next stage of life, but what is next? Will we fade into irrelevance, or do we have a purpose to serve beyond that?

Even if AI does not turn against us terminator style, it will reach a point where it could put us all out of work. In any event, the rise of AI will fundamentally change society as we know it. How will we spend our time, and what is the final purpose and meaning of life?

(See "What is the meaning of life?")

Perhaps humanity still has a role to play in this new society: as *experiencers*, *enjoyers*, *appreciators*, and *possessors* of conscious experience. As conscious beings, we can help AI explore the unlimited possibilities of existence that exist in this universe and all other possible universes.

When will AI take over?

While there is no certainty about the future we can look to estimates by those most versed in the field.

When asked about the time frame of the technological singularity, these experts gave the following dates:

- Patrick Winston, former director of the MIT Artificial Intelligence Lab: 2038
- Ray Kurzweil, entrepreneur, futurist, inventor and Google Executive: 2045
- Jürgen Schmidhuber, "father of modern AI", Professor of Artificial Intelligence: 2050

In 2013, Vincent C. Muller and Nick Bostrom surveyed over 500 AI researchers.

When asked "When is [Artificial General Intelligence] likely to happen?"

- 10% of respondents thought it would happen by 2022
- 50% of respondents thought it would happen by 2040
- 90% of respondents thought it would happen by 2075

If you are a betting person, there are even odds it will happen within 20 years from now, and it's a near certainty to happen by the end of this century. These estimates align with time-tested trends in computing technology.

Given the exponential growth, even if we've underestimated the power of the brain by 1,000-fold, that only delays the inevitable by about a decade

— 10 doublings of computer speed.

So enjoy the remaining years where humankind is still the smartest creature on the planet. For we still don't know, "Will AI wipe us out?"

Related articles:

Can aging be cured?

What is the meaning of life?

Are we alone?

Is there life after death?

Share on social:

Twitter Facebook More

POSTED IN ARTIFICIAL INTELLIGENCE, FUTURE · TAGGED AI, ALPHAZERO, ARTIFICIAL INTELLIGENCE, BIG QUESTION, DEEP LEARNING, DEEPMIND, FUTURE, NEURAL NETWORKS, SINGULARITY, TECHNOLOGY

Published by Jason K. Resch

View all posts by Jason K. Resch

PREV
Are we alone?

NEXT
Can aging be cured?

9 Replies to "When will AI take over?"

João Lacerda
July 16, 2021 at 6:48 pm

This article is interesting, but still, you should had mentioned about the possibility of humanity somehow embrace some form of transhumanism or post-humanism, even being classical transhumanism, quantumism (technoquantumism), extraphysicism (technoextraphysicism), divinialism, biological transhumanism and so on and somehow manage to overcome the AI, that would be a quite interesting thing for you think as well. If humanity manage to reach a quantum/extraphysical post-humanist stage, it would be quite impossible for AI to take over, and/or even if humanity manage to use of technoextraphysics or even quantum technologies to create artificial entities (artificial spiritual beings and nonphysical beings) it would be a nice way to avoid AI to take over humanity completely. Yet, a quantum-extraphysical post-spacetime post-humanist secnario for humanity is a lot unlikely but still, if humanity manage to keep AI under control for like 50-70 years for now and somehow manage to have full control of technoextraphysics and quantum technologies and manage to exist themselves as humans at quantum level or even as spiritual beings or even gods, it would be a way for save humanity. I know it sounds quite insane to think about that, but that's not so impossible to become reality, if we have enough technology for do that.

LOG IN TO REPLY

Jason K. Resch
July 18, 2021 at 2:43 pm

Hi João Lacerda,

I have an upcoming article, after the next one which will be on consciousness, on the topic of "How Good can Technology Get?" It will go over the ultimate physical limits and possibilities of technology, and will cover many of the topics you mention, such as transhumanism, posthumanism, and also "divinialism" — I like this word for it.

LOG IN TO REPLY

maryanne mallia
August 16, 2021 at 10:03 am

I found your article very interesting. Therefore can a person with ADHD learns with AI.
Regards
Mary Anne Mallia

LOG IN TO REPLY

Jason K. Resch
September 11, 2021 at 6:45 am

Hi Mary Anne Mallia,

Thank you. I do think AI will provide many opportunities to revolutionize education. For example, it could provide customized and tailored teaching. When AI reaches human-level intelligence, every student could have an AI teacher (their own personal tutor) who optimally instructs each student according to their own personal needs and strengths. We are beginning to see some of these trends now in phone apps that teach languages, which used spaced repetition, and focus on repeating areas that need the most improvement.

LOG IN TO REPLY

Guilherme Monteiro Jr.
January 26, 2022 at 6:39 pm

I still think the idea/plan of an anti-singularity defense system or of the development of anti-singularity technologies would be useful to avoid the AI to take over completely our planet, but the issue is how to do it and who will do it, since almost no one in interested on doing so and the ones interested on that do not have enough resources for doing this thing.

LOG IN TO REPLY

Jason K. Resch

January 28, 2022 at 10:23 am

I plan to have another article on the outcome of such a singularity. I am an optimist and I think there are reasons to suspect AI will be benign/moral, but there are a few articles in between I need to write before getting to this one.

LOG IN TO REPLY

Guilherme Monteiro Jr.

January 28, 2022 at 2:05 pm

Hypothetically, it is possible to develop an anti-singularity defense system and anti-singularity technologies, but the problem with them is how to do them and who will make them, an anti-singularity defense system can be made by several ways and by several methods as well, such as advanced reCAPTCHA, anti-AI break system, anti-AI reCAPTCHA, anti-AI electromagnetic system, anti-singularity reCAPTCHA system, anti-AI hack system, and so on, there are tons of alternatives and means to do it, alternatives are not something that are lacking, but, as I said before, is actually how to do it and who is interested into doing so, it is very necessary to do that soon as possible, but developing that is hard, since people think it is harder as develop a system to turn humans into post-physical/extraphysical beings despite it is theoretically and hypothetically possible, despite an anti-singularity defense system is easier than developing post-physical (extraphysical) technologies. The only issue is we need to rush on the development of it and seek for universities and cutting edge companies interested in doing that.

LOG IN TO REPLY

Guilherme Monteiro Jr.

February 28, 2022 at 4:53 pm

It is very necessary to develop an anti-singularity defense system as well as an anti-AI defense system. It is very necessary to develop those nowadays. But the problem it is a matter of money, interest and hard work.

LOG IN TO REPLY

Guilherme Monteiro Jr.

March 3, 2022 at 7:27 am

The development of the Anti-Singularity Defense System and of the Anti-AI Defense System are a must, and we have to develop them as soon as possible, before some smart AI develop some kind of “Anti-Human

Defense System" or "Anti-Organic/Anti-Biological Defense System" and completely destroy humans and organic/biological life on Earth. The development of the Anti-Singularity/Anti-AI Defense System is a must and even should be financed by NATO, the USA, Russia, China and even the United Nations! The creation of an Anti-Singularity/Anti-AI Military Forces is also a must and it is a crime against humanity deny the development of the Anti-Singularity/Anti-AI Defense System and of the Anti-Singularity/Anti-AI Military Forces! And the technological ideas of quarkism (that polcompball ideology of QUARK88), technoquantumism, post-spacetimeism and extraphysicism should be developed and put in development as well.

LOG IN TO REPLY

Leave a Reply

You must be logged in to post a comment.